RESEARCH SYNTHESIS / PRINTABLE HTML MANUSCRIPT

TAIJITU: From a Pre-Geometric Two-Sector Mechanism to Galaxy-Scale Rotation-Curve Tests

A consolidated scientific narrative of the TAIJITU programme, from the first spectral and phase-transition diagnostics to the latest parametric mapping experiments. This document separates what has been established, what has only been diagnosed, and what remains to be demonstrated.

AUTHOR: JANOS GABOR MELEGH

Abstract

TAIJITU is a two-sector pre-geometric framework in which the observable ("visible") sector is coupled to a hidden sector, and the induced dark correction is generated through a Schur-complement mechanism. The long-term goal of the programme is not merely to fit galaxy rotation curves, but to show that a physically motivated network-level mechanism can generate an effective dark component, produce galaxy-scale observables with a small set of global parameters, and ultimately remain distinguishable from phenomenological Navarro–Frenk–White (NFW) halo fitting.

The programme progressed through three major stages. First, the early phases established the existence of an ε-driven transition, its first-order character, the attractive and infrared-dominant character of the induced correction, and the necessity of the two-sector architecture. Second, Phase 3b showed that the previously observed "repulsive" dark component was not physical but numerical, originating from an ill-conditioned hidden-sector matrix; a positive-definite stabilisation resolved the sign problem. Third, Phases 5–7.3 addressed astrophysical calibration and galaxy-scale fitting. A near-universal velocity normalisation, $\lambda^* \approx 12$, emerged robustly across the 10-galaxy benchmark and remained stable under later extensions. With both model and target curves normalised, Phase 6 produced the first meaningful shape-based result: median shape $\chi^2 = 1.90$, 8/10 galaxies with $\chi^2 < 5$, and 4/10 galaxies in which the TAIJITU dark component improved the fit over the baryon-only reference.

Phase 7 delivered an important negative result: density-dependent coupling and multi-seed averaging made the effective dark profile smoother and more centrally concentrated, but did not improve any fit. The optimiser simply drove the dark contribution to zero, showing that the bottleneck was the observable mapping from Σ_{DM} to $v_{DM}(r)$, not the network generator itself. Phase 7.1 then provided the key diagnostic breakthrough: a parametric monotonic mapping activated the dark sector immediately on a 6-galaxy stress test, while the original and spectral mappings remained effectively "dark-sector off." Phase 7.2b extended this to the full 10-galaxy sample: the dark sector became active in 9/10 galaxies, χ^2 median dropped to 0.699, and RAR RMS improved from 0.346 dex to 0.3075 dex. Phase 7.3 stabilised this trend, achieving ε off-edge in 10/10 galaxies and χ^2 median 0.651, but it did not convert the remaining five difficult galaxies into net wins, and RAR RMS stayed above the target 0.25 dex.

The main scientific conclusion is therefore mixed but strong. The programme has already established the internal two-sector mechanism, its transition structure, the

numerical resolution of the sign problem, the emergence of a near-universal velocity scale, and the fact that galaxy-scale failure in earlier phases was primarily a mapping problem. What has *not* yet been established is that the physically derived TAIJITU mapping, without explicit NFW-like regularisation, can outperform baryon-only baselines or conventional halo templates on the full sample. The immediate research frontier is thus clear: replace the diagnostic parametric taper by a physically motivated observable map derived more directly from the Schur-complement structure, and validate the framework against full observed radial data rather than synthetic Freeman+NFW reference curves.

1. Starting Point, Scientific Aim, and Strategy

TAIJITU began from a conceptual question: can a hidden pre-geometric sector generate an effective dark correction through relational structure alone, without inserting a dark halo by hand?

The initial ambition was larger than curve fitting. The goal was to build a framework with three layers of credibility:

1. **Internal mechanism:** a two-sector construction must generate a mathematically well-defined dark correction with testable scaling structure.
2. **Transition structure:** the coupling parameter ε should not merely tune amplitudes continuously but reveal a genuine regime change in the coupled system.
3. **Astrophysical observables:** the resulting effective correction should produce realistic galaxy-scale velocities with a limited and interpretable normalisation structure.

Accordingly, the research programme was organised in stages. Early phases studied spectral structure, universality, and order parameters. Intermediate phases diagnosed numerical pathologies and fixed them. Later phases moved to galaxy-scale calibration, shape-based fitting, and mapping ablations. This sequencing matters: the strongest claims of the project are not all astrophysical, and the astrophysical claims do not all stand or fall together.

2. The TAIJITU Framework in One Page

At the core of TAIJITU is a block-structured relational operator in which a visible sector is coupled to a hidden sector. Eliminating the hidden sector yields an effective correction in

the visible sector through a Schur complement. In the notation used throughout the project, the induced contribution is denoted Σ_{DM}.

The visible question is then how to turn this induced operator into a radial contribution to galaxy rotation curves. That mapping problem eventually became the central scientific bottleneck of Phases 6–7.3. Crucially, this means one must distinguish between:

- the existence and internal properties of the induced correction,
- the numerical stability of its computation, and
- the astrophysical observable map used to turn it into $v_{DM}(r)$.

Three distinct scientific questions

Q1. Does the two-sector mechanism generate an attractive, structured correction?

Q2. Is the system organised by a real coupling-driven transition rather than by arbitrary tuning?

Q3. Does the observable mapping produce galaxy-scale curves that remain competitive once both amplitude and shape are measured meaningfully?

3. What Was Established Before Galaxy Fitting

3.1. An ε-driven transition exists and is not a numerical mirage

Across the early spectral programme, the coupled system consistently exhibited a critical region near $\varepsilon \approx 0.88$. Finite-size scaling, multiple seeds, and Binder/hysteresis diagnostics supported a genuine transition rather than a smooth crossover. In the project synthesis, the first-order character was summarised by a Binder cumulant minimum $U_4 \approx -0.78$, while the finite-size programme consistently returned a critical region near $\varepsilon^* \approx 0.88$. A power-law fit to the spectral gap gave a critical exponent $\nu \approx 0.112 \pm 0.006$, with the fitted transition location slightly below the derivative-based estimate, as is common in finite-range numerical fits.

ε* ≈ 0.88 CRITICAL REGION	**U4 ≈ −0.78** FIRST-ORDER INDICATOR
ν = 0.112 ± 0.006 SPECTRAL FIT	**5 seeds / 5 sizes** FINITE-SIZE SUPPORT

3.2. The induced correction is attractive and infrared-dominant

The early phase synthesis also established that Σ_{DM} is not an arbitrary residual. It is attractive, scales approximately as ε^2, and is infrared-dominant across tested topologies. This point matters conceptually: the hidden sector does not merely inject random power into the visible sector, but generates a long-scale correction with a stable sign structure.

3.3. The two-sector architecture is necessary

One of the strongest early claims of the programme is architectural, not phenomenological. A one-sector visible-only construction failed to reproduce the same transition structure and effective dark phenomenology. In the internal synthesis, this was summarised by the statement that the relevant response parameter γ shifted from 0 to 1.77 only in the coupled two-sector setting. The consequence is important: in TAIJITU, the hidden sector is not an optional embellishment but the mechanism itself.

Established by the early programme

- A coupling-driven transition exists near ε ≈ 0.88.
- The transition behaves as first order in Binder/hysteresis diagnostics.
- The induced correction is attractive and IR-dominant.
- The two-sector architecture is required to obtain the effect.

4. Phase 3b: The “Repulsive Dark Matter” Problem Was Numerical, Not Physical

Before galaxy-level claims could be taken seriously, the programme had to resolve a major pathology: in earlier computations the induced dark correction sometimes appeared repulsive. Phase 3b showed decisively that this was not new physics but

numerical instability. The hidden-sector matrix Ω_L had a condition number of order 10^{16}, near the limit of double-precision arithmetic. Although Ω_L was formally positive definite, its numerical inverse developed two negative eigenvalues, which in turn contaminated the sign of Σ_{DM}.

PHASE 3B DIAGNOSTIC	VALUE	INTERPRETATION
Condition number of Ω_L	$\sim 10^{16}$	Numerically near-singular in 64-bit arithmetic
Negative eigenvalues of $G_L = \Omega_L^{-1}$	2	Sign corruption introduced by inversion
Stabilised $\mathrm{Tr}(\Sigma_{DM})$	+18.0	Recovered attractive sign after regularisation
Positive diagonal after stabilisation	100%	Sign problem resolved
Median χ^2/dof after stabilisation	0.077	Good qualitative reference-curve agreement

This resolution was scientifically crucial because it removed an entire class of false interpretations. From this point on, the sign of the induced correction could be discussed as physics again. Phase 3b therefore belongs among the strongest established results of the programme.

5. Phase 4: Why an Apparent “Success” Was a False Positive

Phase 4 is important precisely because it was *not* a final success. The shape χ^2 looked excellent on paper: median $\chi^2 = 0.168$ and 9/10 galaxies had $\chi^2 < 2$. But the same phase also showed that the model velocity scale overshot the observed scale by a factor of about 11.25 on average, while λ was frozen at the lower boundary value $\lambda = 1$ for every galaxy. In other words, the fit was not discovering a universal physical normalisation. It was simply never testing that degree of freedom meaningfully.

PHASE 4 QUANTITY	VALUE	SCIENTIFIC READING
Median χ^2	0.168	Apparently good shape agreement
Galaxies with $\chi^2 < 2$	9/10	Looked strong at first sight
Mean v_{pred}/v_{obs}	11.25×	Severe amplitude mismatch
λ	1 for all galaxies	Grid-edge lock, not universality
κ spread	2395×	Large compensating freedom

Phase 4 therefore clarified an essential lesson: amplitude calibration and shape agreement had to be separated. This led directly to the two-step normalisation strategy of Phases 5 and 6.

6. Phase 5: Emergence of a Near-Universal Velocity Scale

Phase 5 provided one of the most important positive results of the astrophysical part of the programme. Once the amplitude was treated explicitly, a near-universal velocity normalisation emerged:

11.7 ± 2.6
MEAN Λ*

2.8×
Λ* SPREAD

22%
Λ* COEFFICIENT OF VARIATION

0.925 ± 0.078
VPRED / VOBS

For the first time, the model's flat-velocity predictions lay on the observed scale after division by a single galaxy-level factor λ*, and the distribution of λ* itself stayed within a relatively narrow band across the 10-galaxy benchmark. The BTFR slope also improved to 1.169, close to the observed 1.224 used in the project. This was a real breakthrough.

At the same time, Phase 5 also diagnosed why χ^2 remained very large (≈178): the target curve itself had not yet been normalised. Thus Phase 5 answered two questions separately: the physical velocity scale *did* work, but curve-shape agreement could not yet be meaningfully assessed.

7. Phase 6: The First Meaningful Shape-Based Galaxy Test

Phase 6 normalised both the model curve and the target curve and thereby produced the first genuinely interpretable shape χ^2. This is the phase at which galaxy fitting became physically meaningful rather than merely diagnostic.

PHASE 6 BENCHMARK	VALUE
Median shape χ^2	1.90
Galaxies with $\chi^2 < 2$	5/10
Galaxies with $\chi^2 < 5$	8/10
Mean λ^*	11.7
λ^* spread	2.8×
Mean v_{pred}/v_{obs}	0.925
Galaxies improved by DM over baryon-only	4/10
BTFR slope	1.17
RAR RMS	0.346 dex

This phase established three things simultaneously. First, the λ^* result was not a one-off accident. Second, the baryon-plus-induced-dark framework could already produce realistic galaxy-scale shapes in many cases. Third, the dark correction was not yet consistently helpful: in 6/10 galaxies it made the fit worse, and both ε and κ sat at the lower grid edge. The problem was therefore no longer amplitude but profile shape.

Interpretation of Phase 6

The model had crossed an important threshold: it could now be judged on shape, not only on scale. But the optimiser preferred the smallest possible dark contribution, which meant that the mapping from Σ_{DM} to $v_{DM}(r)$ was still suppressing the physical role of the hidden sector.

8. Phase 7: A Negative Result with High Scientific Value

Phase 7 introduced three substantial modifications: density-dependent coupling, multi-seed averaging, and an extended sample of 20 galaxies. Qualitatively, these changes improved the raw Σ_{DM} profile: it became smoother and more centrally concentrated. Quantitatively, however, the galaxy fits did not improve at all on the original 10-galaxy benchmark. Every galaxy returned the same shape χ^2 as before, while the optimiser simply moved to the new lower corner of the enlarged grid: $\varepsilon = 0.01$, $\kappa = 0.01$ for all 10 galaxies.

1.90 → 1.90 MEDIAN X² UNCHANGED	**4/10 → 4/10** DM IMPROVEMENT UNCHANGED
0/10 GALAXIES IMPROVED VS PHASE 6	**ε = κ = 0.01** NEW LOWER-GRID CORNER

This was not a failure of the programme but a narrowing of the diagnosis. The hidden sector and network generator *were* changing the profile, but the observable map still interpreted the dark contribution as harmful and switched it off. The root problem had been localised: the bottleneck was the mapping logic in the curve() step, not the existence of Σ_{DM} itself.

What Phase 7 still contributed positively

The 20-galaxy extension improved the robustness of the velocity normalisation: $\lambda^* = 11.94 \pm 2.20$ with a coefficient of variation of 18.5%, and $v_{pred}/v_{obs} = 0.951 \pm 0.066$. But these gains primarily strengthened the baryonic baseline plus global normalisation; they did not yet validate the dark-sector mapping.

9. Phase 7.1: The Mapping Bottleneck Was Exposed by Ablation

Phase 7.1 compared three observable mappings on a 6-galaxy stress test: the original mapping, a parametric monotonic mapping, and a spectral mapping. The outcome was unambiguous.

MAPPING MODE	MEDIAN X^2	E OFF EDGE	MEAN $V_{DM,MAX}$	READING
Original	2.03	0/6	3×10^{-5}	Dark sector effectively off
Parametric	0.006	6/6	49 km/s	Dark sector activated immediately
Spectral	2.03	0/6	6×10^{-5}	No practical gain in this implementation

This was the central diagnostic breakthrough of the whole late programme. The original mapping and the spectral attempt both left the induced dark term effectively six orders of magnitude below the baryonic contribution. The parametric mapping, by contrast, activated a dark contribution of order 15–105 km/s and moved ε into the 0.10–1.00 range on the stress test.

Yet this came with an important caveat: the parametric mapping is an **ablation proof**, not a final physical proof. It contains an explicit monotonic, NFW-like prior. It therefore demonstrates that the previous failure was indeed a mapping problem, but it does not yet prove that the physical Schur-complement output naturally yields the required halo form without regularisation.

10. Phase 7.2b: First Full-Sample Activation of the Dark Sector

Phase 7.2b completed the full 10-galaxy parametric run and, for the first time, made the dark sector active across nearly the whole benchmark.

1.8548
MEAN X^2

0.6990
MEDIAN X^2

9/10
E OFF EDGE

0.3075 dex
RAR RMS

PHASE 7.2B SUMMARY	VALUE	COMPARISON
Median χ^2	`0.6990`	Improved from 1.90 in Phase 6
Galaxies with $\chi^2 < 1$	`5/10`	Strong subset
Galaxies with $\chi^2 < 2$	`6/10`	Majority
Galaxies with $\chi^2 < 5$	`8/10`	Same broad consistency class
Mean λ^*	`12.3 ± 2.6`	Stable vs earlier phases
v_{pred}/v_{obs}	`0.929 ± 0.074`	Stable
DM improves over baryon-only	`5/10`	Better than 4/10, still below target 6/10
RAR RMS	`0.3075 dex`	Improved from 0.346 dex, still above target
BTFR slope	`1.166`	Close to observed 1.224

This run was scientifically important for two reasons. First, it verified on the full sample what the stress test had already implied: the earlier dark-sector silence was not fundamental, but mapping-induced. Second, it showed that the parametric family was helpful but not yet sufficient. Five galaxies became strong wins, while five remained net losses relative to the baryon-only reference.

The parameter pattern was also revealing. In Phase 7.2b, r_s/R_d sat at 0.10 for every galaxy, and α frequently hit the upper edge of the tested range. This meant that the old "ε and κ sit on the floor" problem had been replaced by a new one: the optimiser now demanded the smallest scale radius and often the steepest inner profile allowed by the parametric family.

11. Phase 7.3: Stabilisation Without Full Conversion

Phase 7.3 expanded the parametric grid to test whether the Phase 7.2b boundary behaviour was only an artefact of too narrow a halo family. The expanded run improved some aggregate quantities, but it did not change the winner/loser structure of the sample.

PHASE 7.3 SUMMARY	VALUE	RELATIVE TO 7.2B
Mean χ^2	`1.7603`	Slightly better
Median χ^2	`0.6508`	Slightly better
ε off edge	`10/10`	Improved from 9/10
DM improves over baryon-only	`5/10`	Unchanged
RAR RMS	`0.3124 dex`	Slightly worse than 7.2b, still better than 6
Mean λ*	`12.2 ± 2.5`	Stable
v_{pred}/v_{obs}	`0.927 ± 0.076`	Stable
BTFR slope	`1.165`	Stable

The new grid succeeded in one clear sense: α no longer sat on its upper edge, and the mean scale-radius fraction moved down to 0.039 with a range 0.01–0.08. But four galaxies still preferred the smallest tested scale radius, and the same five difficult systems remained difficult. The programme had therefore reached a new, narrower diagnosis: the issue was no longer “why is the dark sector switched off?” but “why does the currently allowed halo family fail on a persistent subset of the sample?”

What Phase 7.3 really proved

The mapping obstruction has been broken. The induced dark correction can be made dynamically relevant and beneficial on half the benchmark sample without destroying the λ* result or the BTFR trend. But the present parametric family still does not provide a universal full-sample solution.

12. What Has Been Scientifically

Established

CLAIM	STATUS	WHY THIS IS STRONG
An ε-driven transition exists	ESTABLISHED	Finite-size, seed-averaged, Binder/hysteresis-supported
The transition behaves as first order	ESTABLISHED	Binder and hysteresis diagnostics
Σ_{DM} is attractive and IR-dominant	ESTABLISHED	Repeated topology tests and early synthesis
The two-sector architecture is necessary	ESTABLISHED	Visible-only constructions fail to reproduce the response structure
The earlier repulsive solution was numerical, not physical	ESTABLISHED	Conditioning analysis plus stabilisation resolved the sign
A near-universal velocity normalisation λ* exists on the benchmark	ESTABLISHED	Persistent $\lambda^* \approx 12$ with spread 2.4–2.8× across later phases
The framework yields realistic galaxy-scale amplitudes	ESTABLISHED	v_{pred}/v_{obs} remains near 0.93–0.95
Galaxy-level failure in the original late phases was a mapping bottleneck	ESTABLISHED	Parametric ablation immediately activates the dark sector

13. What Has Not Yet Been Proven

OPEN QUESTION	CURRENT STATUS	WHY IT REMAINS OPEN
A purely physical observable map from Σ_{DM} to $v_{DM}(r)$	OPEN	The successful mapping so far is parametric and contains an explicit NFW-like regularisation
Full-sample superiority over baryon-only baselines	OPEN	Current best runs improve 5/10 galaxies, below the project target of 6/10
RAR scatter at the target level	OPEN	RAR RMS improved from 0.346 to ~0.31 dex, but not below 0.25 dex
Stable finite-size phase-transition diagnostics in the late mapping programme	OPEN	Later spectral tests give a signal, but not yet a settled scaling story
Validation against fully observed radial data	OPEN	Most late benchmark comparisons still use synthetic Freeman+NFW reference curves rather than full SPARC radial likelihoods
Physical interpretation of parametric amplitudes and halo-shape parameters	OPEN	κ remains weakly identified, and halo-family boundaries are still active

14. Scientific Interpretation of the Current State

The central lesson of the entire TAIJITU programme is that **mechanism, numerics, calibration, and observable mapping must be separated**. Once this is done, the landscape becomes much clearer than it first appeared.

The early work succeeded in establishing a real coupled-system mechanism. Phase 3b removed a major numerical artefact. Phase 5 established a persistent global velocity scale. Phase 6 produced the first meaningful shape test. Phase 7 narrowed the late-stage failure to a specific observable-mapping bottleneck. Phase 7.1–7.3 then showed that once this bottleneck is relaxed, the dark sector becomes active and often highly beneficial, but not yet universally so.

In that sense, the programme has not “failed to fit galaxies.” It has reached a sharper and more demanding stage: it now has to convert a successful diagnostic mapping into a physically derived mapping that preserves the dark sector’s newly demonstrated relevance without importing its desired structure by hand.

15. Next Research Directions

15.1. Replace the diagnostic parametric taper by a physically motivated map

The immediate priority is no longer blind parameter expansion. It is to derive a better observable map from Σ_{DM} itself. Two routes are especially motivated:

1. **Spectral low-mode mapping:** use the low-eigenvalue structure of the full Schur complement, not only its diagonal, to define an effective enclosed mass.
2. **Kernel or potential-based mapping:** interpret Σ_{DM} through an effective potential kernel rather than a local shell-density ansatz.

15.2. Expand the halo family only where the data demand it

Phase 7.3 suggests that the difficult galaxies are not random failures. They form a stable loser subset. The most efficient next step is therefore a targeted family expansion on those cases: gNFW with an extra inner-slope degree of freedom, Einasto-type profiles, or profile priors tied to galaxy class rather than a single universal template.

15.3. Move from synthetic targets to full observed radial likelihoods

The benchmark used so far was essential for rapid internal iteration, but the next publication-grade astrophysical test must fit real observed radial data points with their uncertainties. That is where the framework becomes externally testable in the ordinary empirical sense.

15.4. Keep the phase-transition programme alive, but separate it from the galaxy-fitting loop

The late project already showed that rotational-curve-based order parameters are the wrong place to measure the transition once the observable map suppresses the dark sector. Spectral Binder, gap, and finite-size observables should therefore continue on a separate track.

15.5. Preserve the strongest existing result while extending the fit family

One practical rule should govern all future work: do not destroy the λ* result. The emergence of λ* ≈ 12 with moderate spread is one of the most robust outcomes of the entire programme. Any new mapping or profile family that improves χ^2 at the cost of destroying that scale structure would be a trade rather than a genuine advance.

16. Conclusion

TAIJITU has already crossed several thresholds that many exploratory frameworks never reach. It has identified a two-sector mechanism, exhibited transition structure, survived a major numerical pathology, produced a stable galaxy-scale velocity normalisation, and isolated the precise step at which astrophysical performance is now limited.

The strongest fair claim at the present stage is this:

Core conclusion

TAIJITU has **scientifically demonstrated** a viable two-sector pre-geometric mechanism with an attractive induced correction, a real coupling-driven transition, a resolved numerical sign problem, and a robust near-universal velocity scale. It has also demonstrated that the main remaining obstacle to full galaxy-scale success is the observable mapping from the induced operator to the effective halo contribution.

The strongest fair caution is equally clear:

Core caution

The current parametric success is a **diagnostic success, not yet a final physical proof**. The decisive next milestone is a physically motivated map that reproduces the Phase 7.1–7.3 dark-sector activation without relying on an explicit NFW-like taper.

If that milestone is reached—and if the same structure survives full radial data fitting—the TAIJITU programme will move from an internally coherent research framework to a quantitatively competitive astrophysical theory candidate.

Appendix A. Phase-by-Phase Numerical Summary

PHASE	MAIN OUTCOME	REPRESENTATIVE NUMBERS	STATUS
1–2	Transition and mechanism	$\varepsilon^* \approx 0.88$; Binder U4 ≈ -0.78; attractive, IR-dominant Σ_{DM}; two-sector necessity	ESTABLISHED
3b	Repulsive artefact resolved	$\mathrm{cond}(\Omega_L) \approx 10^{16}$; 2 negative eigenvalues in G_L; $\mathrm{Tr}(\Sigma_{DM}) \to +18.0$ after stabilisation	ESTABLISHED
4	False-positive low χ^2 diagnosed	median $\chi^2 = 0.168$, but $v_{pred}/v_{obs} \approx 11.25\times$ and $\lambda = 1$ for all galaxies	DIAGNOSTIC
5	Velocity-scale breakthrough	$\lambda^* = 11.7 \pm 2.6$; spread 2.8×; $v_{pred}/v_{obs} = 0.925 \pm 0.078$; BTFR slope = 1.169	ESTABLISHED
6	First meaningful shape test	median $\chi^2 = 1.90$; 8/10 with $\chi^2 < 5$; DM improves 4/10; RAR RMS = 0.346 dex	ESTABLISHED
7	Negative result localises bottleneck	no improvement vs Phase 6; $\varepsilon = \kappa = 0.01$ for all 10; λ^* on 20 galaxies = 11.94 ± 2.20, CV = 18.5%	HIGHLY INFORMATIVE NEGATIVE RESULT
7.1	Mapping ablation breakthrough	6-galaxy test: original median $\chi^2 = 2.03$, parametric = 0.006, spectral = 2.03; ε off-edge 6/6 only in parametric mode	DIAGNOSTIC BREAKTHROUGH
7.2b	First full-sample parametric activation	median $\chi^2 = 0.699$; ε off-edge = 9/10; DM improves 5/10; RAR RMS = 0.3075 dex	PARTIAL SUCCESS
7.3	Stabilisation with expanded grid	median $\chi^2 = 0.651$; ε off-edge = 10/10; DM improves 5/10; RAR RMS = 0.3124 dex	PARTIAL SUCCESS

Authoring note. This HTML manuscript is designed for printing, annotation, and further refinement into a formal paper or preprint.
The latest galaxy-scale runs (Phases 7.2b and 7.3) are interpreted conservatively as partial successes and mapping diagnostics, not as final astrophysical proof.